ARMY RESEARCH LABORATORY

Initial Empirical Research with an Experimental Secure Web Portal of Electronics Records Archives

by Binh Nguyen and Glenn Racine

ARL-TR-3676 **November 2005**

NOTICES

Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.

Army Research Laboratory
Adelphi, MD 20783-1197

ARL-TR-3676 **November 2005**

Initial Empirical Research with an Experimental Secure Web Portal of Electronics Records Archives

Binh Nguyen and Glenn Racine
Computational and Information Sciences Directorate, ARL

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED (From - To)
November 2005	Final	

4. TITLE AND SUBTITLE
Initial Empirical Research with an Experimental Secure Web Portal of Electronics Records Archives

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Binh Nguyen and Glenn Racine

5d. PROJECT NUMBER
4FVAV1

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
U.S. Army Research Laboratory
ATTN: AMSRD-ARL-CI-CN
2800 Powder Mill Road
Adelphi, MD 20783-1197

8. PERFORMING ORGANIZATION REPORT NUMBER
ARL-TR-3676

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
U.S. National Archives and Records Administration
8601 Adelphi Road
College Park, MD 20740-6001

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT
This document reports the results of an initial empirical research with an experimental secured portal of sensitive electronic records archives. The experimentation focused on measuring the effective data-transfer rates (throughputs) of the portal under various security settings of the deployed security products. The measured throughputs were preliminary and for internal uses only; they did not indicate or predict the actual capability and the capacity of the portal environment. This work was part of the ARL cooperative agreement number DAAD19-03-2-0018, an enabling vehicle for a joint science-and-engineering research project of ARL and Georgia Tech Research Institute (GTRI). The purpose of this collaborative work was to facilitate the processing and the protection of distributed authentic electronic records archives (ERA) for the U.S. National Archives and Records Administration (NARA).

15. SUBJECT TERMS
Secured web portal, overhead measurements

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Binh Nguyen
Unclassified	Unclassified	Unclassified	UL	34	19b. TELEPHONE NUMBER (*Include area code*) (301) 394-1781

Standard Form 298 (Rev. 8/98)
Prescribed by ANSI Std. Z39.18

Contents

List of Figures

List of Tables

Preface

The U.S. Army Research Laboratory (ARL) has prepared this technical report for submission to the U.S. National Archives and Records Administration (NARA). ARL conducts basic and applied research to provide the technological competitive edge for the U.S. Army. NARA is the record-keeper of the Nation; it is the steward of irreplaceable electronic and non-electronic collections documenting our Nation's experience, the actions of government, and the rights and entitlements of our citizens.

This document reports the findings of initial hands-on research with an experimental secured web portal designed to connect and operate in the Internet. The main purpose of the portal is to provide a relatively low-cost, convenient, secured, and effective means for collaborative processing of sensitive electronic records archives (ERA). The portal is envisioned to be a centralized repository of raw and processed ERA, as well as tools and documents. The portal is being developed under the Presidential Electronic Records Pilot Operating System (PERPOS) program, a NARA-supported research program being performed by the Georgia Tech Research Institute (GTRI), and hence, it is called the PERPOS portal. The portal and its computing and communications equipment are being built as a high-performance test bed capable of providing a suitable environment for conducting empirical research to find practical solutions for protecting and securing sensitive ERA in public networks.

INTENTIONALLY LEFT BLANK.

Executive Summary

This document reports the results of a hands-on interaction with an experimental secure web portal of sensitive electronic records archives. The task focuses on (1) providing preventive services that block unauthorized access to the secured portal computing environment and (2) measuring the effective data-transfer rates of the portal under various settings of the deployed security products. Experimentation with the secured portal revealed its capability and compliance with federal information processing standards (FIPS), issued by the National Institute of Standards and Technology (NIST):

- The portal could provide authentication and integrity services using two versions of the secured socket layer protocols (SSLv2 and SSLv3). The SSLv2 is obsolete, and the SSLv3 is not approved by the government. The SSLv3 and the first version of the transport layer security protocols (TLSv1) are functionally equivalent, but not interoperable. The TLSv1 protocols are the only government-approved protocols for securing sensitive electronic data.
- The portal could offer confidentiality services using only single and triple data encryption standards called DES and Triple-DES, respectively. These standards are government-approved symmetric data encryption algorithms, but they are considered obsolete and being replaced by the advanced encryption standard (AES) algorithms.
- The portal test bed employed a mix of Gigabit Ethernet (1000 Base-T) and Fast Ethernet (100 Base-T) network devices. Some of the network hubs, switches, cables, and network interface cards were designed for a Fast Ethernet network.

The measured data-transfer rates were considered preliminary and for internal use only. They did not indicate or predict the actual capability and the capacity of the portal environment. These initial results provided encouraging evidence that forms the basis for the coming efforts; therefore, ARL recommends the following actions for the portal:

- Acquire and use government-approved cryptographic modules implementing the TLS protocols and the AES symmetric encryption algorithms. Successful integration of these modules will enable the portal to offer the highest protection to sensitive electronic presidential records in a public network.
- Modernize the hardware used in the test bed. Hardware includes workstation-class notebook and desktop computers, network switches, routers, cables, and network interface cards capable of supporting Gigabit Ethernet network technologies.
- Continue conducting empirical experiments evaluating the performance overhead induced by the deployment of defensive IA products at the portal.

- Evaluate open-source, secured, virtual private network products that use cryptographic tunneling protocols to provide authentication, confidentiality, and integrity services for sensitive electronic records archives in public networks.
- Continue assessing other types of government-validated information-assurance products suitable for the protection of the portal. The target IA products include intrusion detection technologies and products, antiviral software, and security management tools.

1. Introduction

1.1 Background

This work is part of the U.S. Army Research Laboratory (ARL) cooperative agreement number DAAD19-03-2-0018, Modification 002, July 1, 2004. The agreement enables ARL to perform a joint science-and-engineering research project with the Georgia Tech Research Institute (GTRI), Atlanta, GA. This collaborative work facilitates the distributed processing and the protection of distributed electronic records archives (ERA) for the National Archives and Records Administration (NARA). The agreement calls for the execution of two main tasks: (1) automated content analysis and information extraction, and (2) information assurance (IA) for distributed processing of ERA. GTRI independently performs and individually reports the progress of the first task and collaborates with ARL to build a high-performance, secure web portal of sensitive ERA. Researchers at GTRI concentrate their efforts on the functional aspects of the portal under the Presidential Electronic Records Pilot Operating System (PERPOS) project. Researchers at ARL focus on the protection of the portal and its contents. The IA task searches for the best IA practices and recommends technological products that can be applied to the problem of processing, securing, and protecting irreplaceable sensitive ERA in the *Internet2* environment.

1.2 Scope

This document reports the very first results of hands-on interaction with an experimental secure web portal at GTRI during the reporting period. The task focuses on (1) providing preventive services that block unauthorized access to the secured portal computing environment and (2) evaluating the performance of the portal under various settings of the deployed security products. The main purposes of this document:

- Reporting work activities and accomplishments during the reporting period
- Reporting encountered technical barriers and strategies for overcoming them
- Recommending research activities to be carried out during the forthcoming phase

The intended audience of this report includes ARL and NARA administrators and managers, ERA and IA researchers, and information technology personnel.

The next section describes specific planned tasks, reports the status of each task, and explains the method by which each task was accomplished. Section 3 concludes the report and recommends research activities to be accomplished during the next phase. Section 4 acknowledges the contributions of other team members. Section 5 lists commonly used acronyms. Section 6 includes information about products used in the project.

2. Tasks

The assigned IA tasks to be performed during the reporting period include two subtasks: (1) experimenting with the deployed IA products and (2) measuring the performance costs of the deployed security products. The test bed environment will be first delineated, and then the descriptions and the status of each task and the methods by which the tasks were carried out are reported in this section.

2.1 Test Bed Environment of the Portal

The initial configuration of the test bed implements the defense-in-depth strategy as depicted in figure 1. The test bed is a private network being protected by layers of defense. The outer layers execute the overarching security policies imposed by Georgia Institute of Technology (*Gatech)* and GTRI. These policies are inviolable, and therefore, they are beyond the reach of the PERPOS project. The portal and the private subnet apply the local security policies that were designed and controlled by the researchers of the PERPOS project. These policies themselves are experimental and subject to incremental changes in response to the needs of the project, while still complying with the overarching policies.

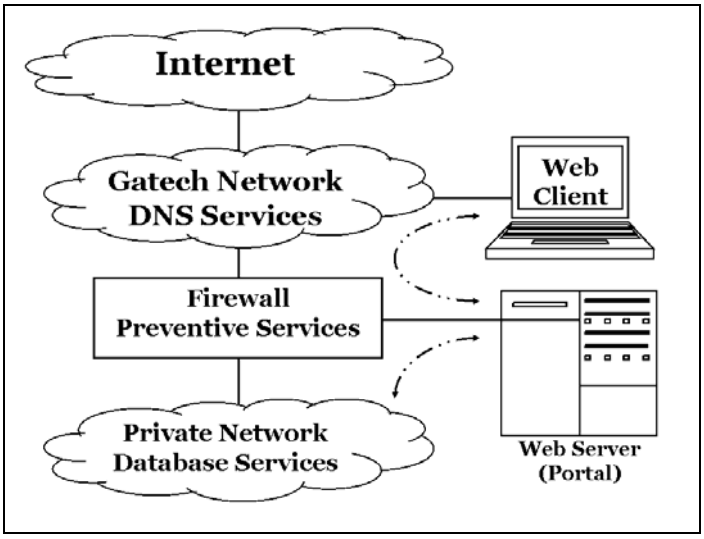

Figure 1. Experimental test bed environment.

The test bed is protected by a firewall that controls incoming and outgoing data traffic and creates a buffer zone between the portal computing environment and other untrusted networks, including the Internet and other internal *Gatech* subnets. The buffer zone is also known as the demilitarized zone (DMZ), the innermost perimeter network defense of the portal test bed. This zone is basically a subnet connected to a network interface of the deployed firewall. To increase reliability of the PERPOS portal system, the deployed computing hardware all used a data

storage technology called “redundant array of independent disks level 1” (RAID1) to create an exact copy of all of data on two disks.

During the experimentation, the DMZ had one computer, a Dell PowerEdge 1750, equipped with dual Intel Xeon processors running at 3.02 GHz clock speed, two Gigabytes of PC2100 memory, two 73-Megabytes hard drives in a RAID1 mirror, and dual Gigabit Ethernet network interface cards. This computer runs the Red Hat Enterprise Linux 3 Advanced Server operating system (OS) and hosts two hypertext transfer protocol (HTTP) servers: an Apache web server and an Oracle 10g application server. Each server runs on a different communication port number. Each communication session is associated with a unique port number. The port numbers are subject to change at any time at the discretion of GTRI. Throughout this document, the term “web server” refers to the Apache web server, and the term “portal” refers to the Oracle application server.

Client computers used in the experimentation included a remote computer and a local computer. The remote computer, being located about ¾ mile away from the portal, was a desktop computer running the Microsoft Windows operating system and *Cygwin* emulation software that offers Unix/Linux functionality in Microsoft Windows environments. It was connected to the *Gatech* network through its 10/100 Mbps Ethernet network interface card. The local computer was a high-performance computer running the Linux operating system and equipped with a Gigabit Ethernet network interface card. Both clients ran the same downloading scripts that ARL had created to support the evaluation of the PERPOS computing network environment. The scripts were written in the Bash shell command language, readily executable in the *Cygwin* and the Linux environments running the Bash shell command.

The test bed was sometimes configured to form an isolated network operating without domain name system (DNS) services by removing the cable connecting the firewall to external networks. The DNS services were normally provided by the *Gatech* network to translate communicating computer names into network addresses. Without DNS services, all computers in the isolated network test bed used their arbitrarily assigned numerical network addresses for communication purposes.

The initial evaluation and experimentation with the portal test bed required the generation of test data files; therefore, simulated test data files having different sizes were created for the web server and the portal. The creation of the test files were accomplished by running a C program designed and implemented by ARL. The sizes of the files were 0 KB, 1 KB, 10 KB, 100 KB, 1 MB, 10 MB, 100 MB, 500 MB, and 1 GB, where 1 K = 2^{10}, 1 M = 2^{20}, 1 G = 2^{30}, and 1 B = 8 bits. As the focus of the experimentation was on large files, small sized data files (under 10MB) were generated only for the comparison and graphing purposes. Discussions of results refer to large files whose sizes were 10 MB or larger because they typify the estimated sizes of future electronic records archives (GB or larger).

A copy of each test file was stored in the web server as a regular disk file and in the Oracle 10g database as a data item, located in a data field within the Oracle database. Each set of test data files had its own index file containing the hyperlink uniform resource identifiers (URIs) of the test data files. If the URIs referred to the Apache server, then the contents of the required data file were read from a disk file and then sent to the requesting client. On the other hand, if the URIs referred to the portal, then the Oracle database server provided the contents of the requested data file.

To monitor the network traffic in the portal environment, two open-source network-protocol analyzers were employed in the experiments: *tethereal* and *ssldump*. The *tethereal* tool was used for monitoring and sometimes capturing all types of traffic moving within the private network, and the *ssldump* tool was used to monitor the traffic whenever the server was operating in secure mode. The *tethereal* tool was readily available in the computers that ran the Red Hat Linux operating system, but the *ssldump* tool had to be downloaded from the Internet, compiled, and installed in the computer hosting the web server and the portal.

The main focus of the initial experimentation was to gauge the functional behavior of the portal. As the portal, powered by the Oracle 10g application server, was based on the Apache technology, an Apache web server was set up to serve as a proving ground for experimenting with output log formats as well as for developing and testing the Bash shell scripts that were used as downloading test data files.

The downloading scripts were designed and implemented to provide an automated interaction with the web server and the portal for different purposes. These purposes included measuring effective data-transfer rates and also extracting a list of cipher suites that the PERPOS portal can handle using two versions of secure socket layer (SSLv2 and SSLv3) and transport layer security (TLSv1) protocols. Therefore, the Apache server existed to serve three main testing purposes. The first purpose was to verify the intended performance of the shell scripts that were used for downloading the test data files and to ensure the capturing of the associated performance data. The second purpose was to acquire further knowledge of the behavior of a secured and unsecured web server in preparation for dealing with the actual portal. The third purpose was to experiment with web server configurations, output log formats, transport layer security options, and the availability of cipher suites.

A cipher suite is a specification of cryptographic algorithms and protocols used in a secure session. For example, the specification “TLS_RSA_WITH_AES_256_CBC_SHA” (AES256-SHA) is the preferred cipher suite for transferring ERA between the PERPOS portal and a client computer. This specification calls for the use of (1) the TLSv1 protocols, (2) the public-key encryption methods invented by Rivest, Shamir, and Adleman (RSA) for key agreement and authentication, (3) the advanced encryption standard (AES) using 256-bit cryptographic keys operating in cipher-block chaining (CBC) mode for encryption and decryption, and (4) the secure

hash standard (SHA) for data integrity. The OpenSSL project has implemented many cipher suites, which can be found at its web site (http://www.openssl.org/docs/apps/ciphers.html).

Experimentation with the secured PERPOS portal revealed its current capability and its compliance with federal information processing standards (FIPS), issued by the National Institute of Standards and Technology (NIST). The first finding was that the portal could handle only the SSLv2 and SSLv3 protocols. The SSLv2 is obsolete, and the SSLv3 is not approved by the government. The SSLv3 and the TLSv1 ciphers are functionally equivalent, but not interoperable, and the TLSv1 protocols are the only protocols that the government currently approves for securing sensitive data in public networks. The second finding was that the portal could use only single and triple data encryption standards called DES and Triple-DES, respectively. These standards are government-approved symmetric data encryption algorithms, but they are now considered obsolete and being replaced by the AES algorithms.

Therefore, only two cipher suites were available for the experimentation with the portal operating in secured mode using the SSLv3 protocols: the DES-CBC-SHA and the DES-CBC3-SHA cipher suites. Both cipher suites use the RSA methods for key agreement and authentication and the SHA secure hash standard for data integrity. The only difference between the two cipher suites is the use of the symmetric encryption algorithm. The former uses single DES with 56-bit cryptographic keys operating in cipher-block chaining (CBC) mode, and the latter uses triple DES with 168-bit cryptographic keys. Using these cipher suites, the PERPOS portal could provide three basic information assurance services: authentication (RSA), confidentiality (DES), and integrity (SHA) services.

Furthermore, to enable the portal to operate in secured mode and to provide a way for a client to authenticate the PERPOS portal, *Gatech* personnel issued self-signed certificates with 2048-bit keys for the PERPOS portal and for the Apache web server. Whenever a secured session took place, the portal presented its certificates to the requesting client for authentication purposes. Certificates were not generated for clients to simplify initial hands-on experimentation with the portal. The portal thus could not authenticate its requesting clients.

The IA products that were purchased and deployed by GTRI and ARL during the performance period consisted of two government-validated firewalls: a Check Point Firewall-1 Next Generation-Application Intelligence (NG-AI) R55 on a Nokia IP350 appliance and a Symantec Enterprise Firewall with VPN 8.0 for Windows. Both firewalls are sophisticated; they are capable of not only inspecting the headers of individual incoming packets, but also performing many other functions such as keeping track of the state of a communication session and inspecting the payload of inbound packets. For example, they can reject an inbound packet that was not requested by a client or that carries a recognizable malicious signature in its payload. The Checkpoint is a turnkey system equipped with an Intel Pentium III processor running at 700 MHz clock speed and three 10/100 Megabits per second (Mbps) Ethernet network interface cards. The Symantec firewall is a software package executing in a hardened Microsoft Windows

2000 operating system running in a high-performance computer, a Dell PowerEdge 1750 equipped with dual Intel Xeon processors running at 3.06 GHz clock speed, two Gigabytes of PC2100 memory, and four Gigabit Ethernet network interface cards.

An ARL researcher in Atlanta, Georgia, working side-by-side with GTRI researchers, completed the first configuration of the two recently purchased firewalls. Experimentation with their configuration parameters was started to optimize their performance and to evaluate their effects on the performance of the portal and the web server. The results of their evaluation of the firewalls are reported separately in a GTRI technical report entitled "PERPOS Information Assurance" by Jason Kau and Son Nguyen, Atlanta, GA, 31 May 2005.

Early experimentation with the portal measured the effective data-transfer rates of the PERPOS system under various settings of each deployed firewall. The data-transfer rates, also known as effective throughputs, were measured at the client computer. They were computed as the ratio of the size of a downloaded file to the total elapsed time taken for a successful transfer, i.e., *throughput = size of data file/elapsed time*. The elapsed time includes the actual time required to transfer a data file and other overhead time required for setting up, maintaining, and releasing a reliable connection (communication overhead) between the two communicating computers.

The results of these experiments then served as the basis for the configuration and reconfiguration of the target firewall in order to optimize its performance. Besides system configuration settings appropriate for a particular environment, other major factors affecting the performance of the firewall included the type of the processor on which it ran, the amount of, and the type of memory available for running it, and the network interfaces. As the two firewalls deployed in the PERPOS environment ran in dissimilar hardware platforms, comparing their relative performance was not the objective of the experimentation.

2.2 Initial IA Product-Induced Overhead Measurement

Conducting initial performance evaluation and overhead measurement of the portal operating under various security configurations had two main objectives. The first objective was to provide a way for understanding and predicting the performance of the secured portal in a laboratory environment. The second objective was to provide a performance baseline against which subsequent measurements would be compared. Measurements should be made whenever changes are made to the PERPOS portal. Changes that would affect the performance of the portal include hardware and software upgrades, system configurations, or security components and services.

This task called for the performance-overhead evaluation of the portal operating in secured mode by conducting experiments to comparatively measure the overhead associated with deployed security products. The security products included cryptographic algorithms and protocols that enabled the portal to operate in secured mode and the two firewalls that had just been installed and configured. To carry out this task, several test scenarios were designed and conducted to

collect some quantitative data that might provide some initial insights into the performance of the portal under various security configurations.

Each test scenario was conducted under different firewall options and network configurations, but used the same shell script that incorporated a web client tool to download the test data files repeatedly from the portal and from the web server. Two web client tools, the *wget* tool and the *curl* tool, were empirically studied and evaluated. The *wget* tool provides a convenient way for recursive downloading, and it is relatively easier to use. However, it does not support the specification of a cipher suite and a transport security protocol. Because of this, it was used mainly for intermittently downloading a file from the portal or for verifying that the portal was working properly. Whereas, the *curl* tool was incorporated in every downloading script file because it can be specified to use a particular protocol and a specific cipher suite.

The internal working of the downloading scripts operated in three phases. First, it retrieved the contents of a specified index file from the web server or the portal. Second, it parsed the contents to extract the URIs of the linked data files. Each URI specifies the protocol (e.g., HTTP or HTTPS), the name of a web server, and the name of a test data file. If the experiments were conducted in the isolated network environment, the scripts substituted the embedded name of the server with its corresponding numerical address. Third, for each extracted URI, it executed the *curl* web client tool to download a data file 10 times. Each data point plotted in the reported figures is the average of 10 successful downloading of the same file. All downloaded data were written to the "/dev/null" device, effectively discarding them, to save disk storage spaces.

Each execution produced a set of statistical information about a download session. The information included, but was not limited to, the return code (e.g., the code "200" indicates success), the total time (seconds), name lookup time (seconds), the size of the downloaded file (bytes), and the download data-transfer rate (bytes per second). The data were saved in a log file for subsequent analysis using Microsoft Excel to produce graphical charts.

Many tests involving two scenarios were conducted in the PERPOS environment, and some of their results are reported in this document. The first scenario involved a remote client computer. The second scenario involved a local high-performance computer connected to the portal in the same network (one hop). In both scenarios, the portal and the Oracle database management system were running in the same computer.

The first scenario was set up according to figure 2 involving a remote web client named Paladin and the Nokia Checkpoint firewall. The client was located about 3 hops away from the portal. The number of hops was an estimate because many computers in the network disallowed route tracing using the *traceroute* and *tracert* commands. Paladin was a desktop computer equipped with an Intel Pentium 4 microprocessor operating at 1.4 GHz clock speed and a 10/100 Mbps Ethernet interface card.

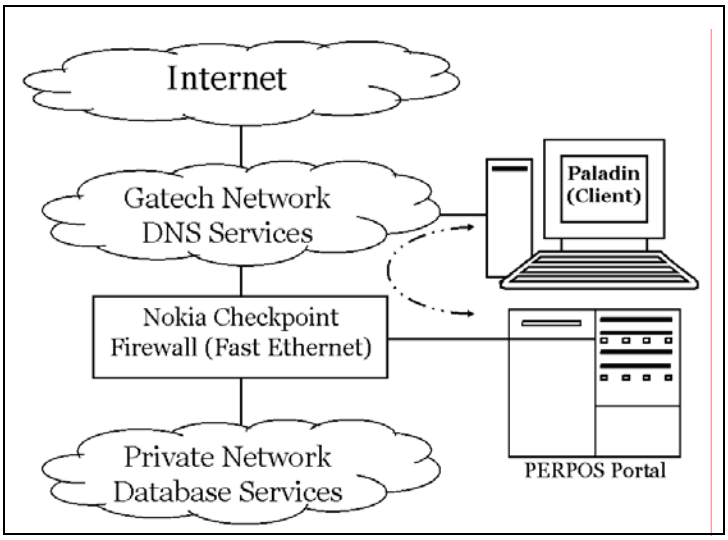

Figure 2. Test scenario 1 (remote connection).

Six tests were conducted in the scenario 1 to measure the performance costs of using security features available at the PERPOS test bed. To realize this objective, effective data transfer rates *(r)* were measured as a function of (1) the operating mode of the portal *(m)*, (2) the cipher suite *(c)* used in a downloading session, and (3) the use of application inspection feature *(a)* of the Nokia Checkpoint firewall, i.e., $r = f(m, c, a)$. Where

$m = \{unsecured, secured\}$

$c = \{null, DES\text{-}CBC\text{-}SHA, DES\text{-}CBC3\text{-}SHA\}$

$a = \{ON, OFF\}$

The following table summarizes the values used in the six tests:

Table 1. Remote retrieval configurations.

Test #	m	c	a	Chart Label
1	unsecured	null	OFF	paladin_unsecure_chkpt-app_insp
2	unsecured	null	ON	paladin_unsecure_chkpt+app_insp
3	secured	DES-CBC-SHA	OFF	paladin_DES-CBC-SHA_chkpt-app_insp
4	secured	DES-CBC-SHA	ON	paladin_DES-CBC-SHA_chkpt+app_insp
5	secured	DES-CBC3-SHA	OFF	paladin_DES-CBC3-SHA_chkpt-app_insp
6	secured	DES-CBC3-SHA	ON	paladin_DES-CBC3-SHA_chkpt+app_insp

Figures 3-7 display the results of these six tests. Test 1 and test 2 assessed the effects of the deployed Checkpoint firewall on the data transfer rates (throughputs) when the portal was operating in unsecured mode. The results of test 1 indicated (1) that the maximum throughputs could reach nearly 80% of the rated speed (100 Mbps) of the Fast Ethernet network interface cards and (2) that the application inspection feature of the Checkpoint firewall reduced about 50% of throughputs (figure 3).

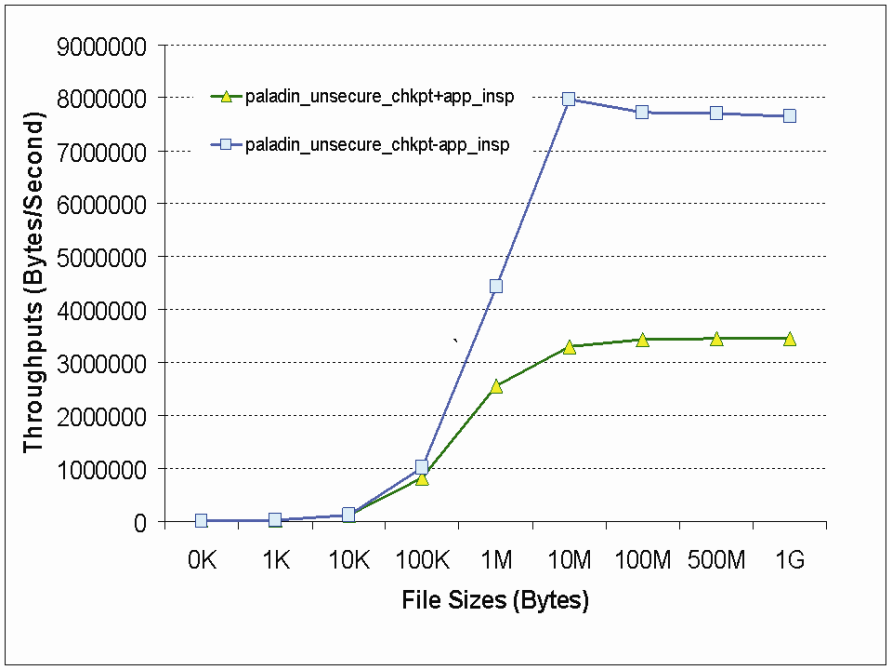

Figure 3. Results of scenario 1: tests 1 and 2.

Test 3 and test 4 evaluated the effects of the deployed Checkpoint firewall on secured data traffic using the cipher suite DES-CBC-SHA and the SSLv3 protocols. The results of these two tests show that the application-inspection feature of the firewall had practically no effects on the secured traffic (figure 4).

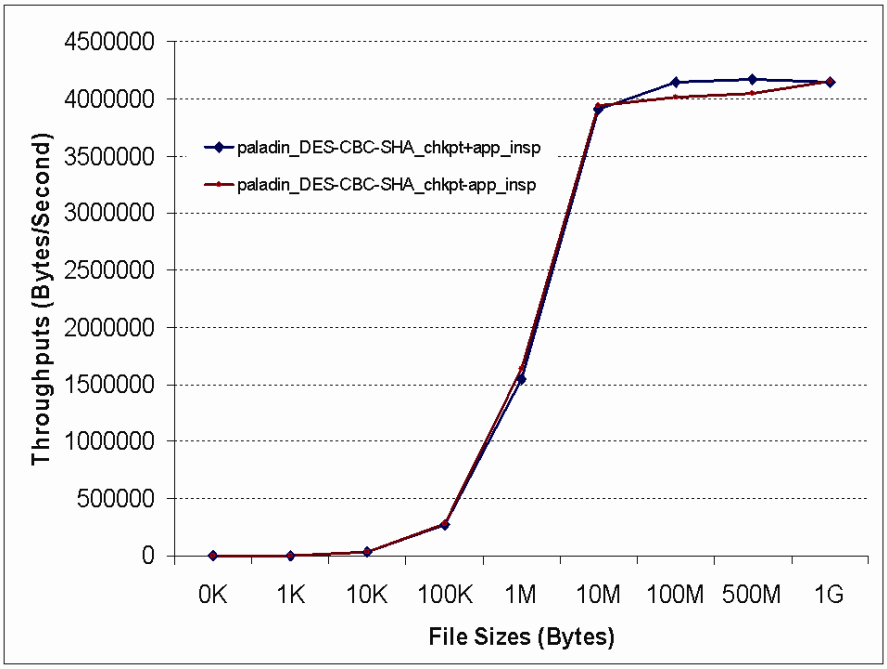

Figure 4. Results of scenario 1: tests 3 and 4.

Test 5 and test 6 evaluated the effects of the deployed Checkpoint firewall on secured data traffic using the cipher suite DES-CBC3-SHA and the SSLv3 protocols. The results of these two tests show that the application-inspection feature of the firewall has a slight impact on the secured traffic (figure 5).

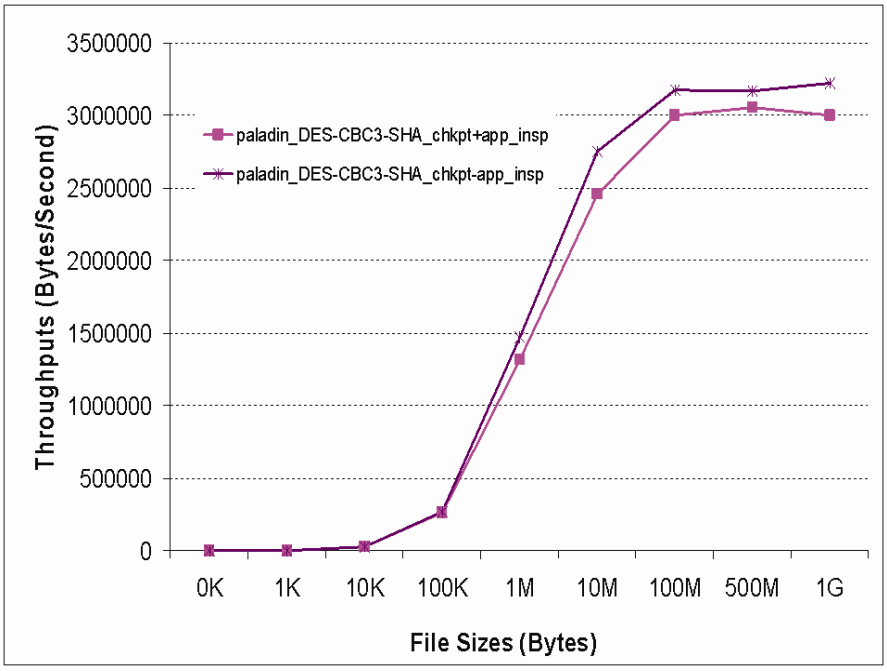

Figure 5. Results of scenario 1: tests 5 and 6.

Figure 6 shows the results of all six tests for comparison purposes. The highest throughputs were achieved when the portal was operating in unsecured mode and the application-inspection feature of the Nokia Firewall was not used. The lowest throughputs were observed when the portal was operating in secured mode using the DES-CBC3-SHA cipher suite and the application-inspection feature of the Nokia Firewall was used.

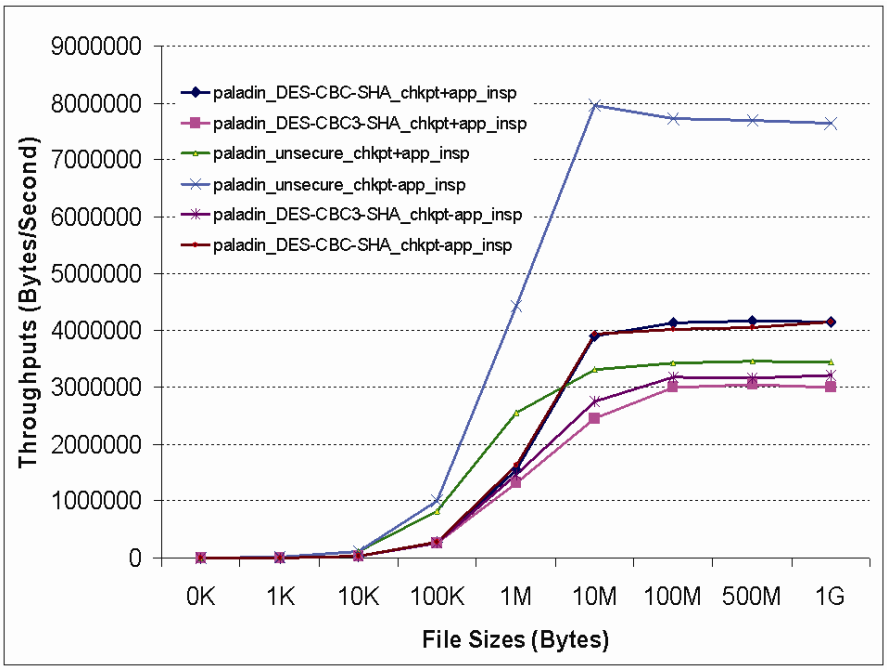

Figure 6. Comparative results of test scenario 1.

Figure 7 shows the performance costs incurred by the Nokia Checkpoint firewall and by various cipher suites using the secured mode of the portal. The use of the DES-CBC3-SHA cipher suite, the strongest cipher suite available at the PERPOS portal, for securely transferring large data files decreased the throughputs to about 40 percent of their maximum. These costs translate into a 250 percent increase of downloading times.

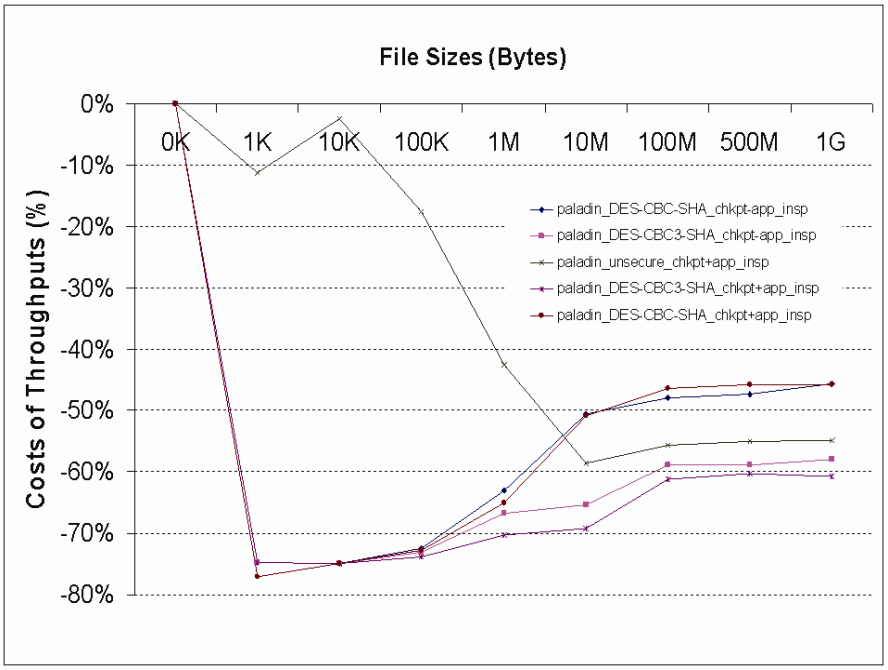

Figure 7. Comparative performance costs of test scenario 1.

The second scenario involved a local web client named Bush41 and the Symantec firewall as shown in figure 8. The client, the server, and the firewall all were high-performance computers equipped with Gigabit Ethernet network interface cards. Category 5e (Cat 5e) network cables were used to connect these computers to form a private network. (The Cat 5e cables are designed for Gigabit Ethernet network.) The client was an HP/Compaq DL380 G3 computer, equipped with dual Intel Xeon processors running at 3.02 GHz clock speed, six Gigabytes of PC2100 memory, and two 36.4-Gigabytes of disk storage configured in a RAID1 mirror, and dual Gigabit Ethernet network interface cards, and the Red Hat Enterprise Linux operating system. The Symantec firewall is a software package executing in a hardened Microsoft Windows 2000 operating system running in a Dell computer model PowerEdge 1750 equipped with dual Intel Xeon processors running at 3.06 GHz clock speed, two Gigabytes of PC2100 memory, and four Gigabit Ethernet network interface cards.

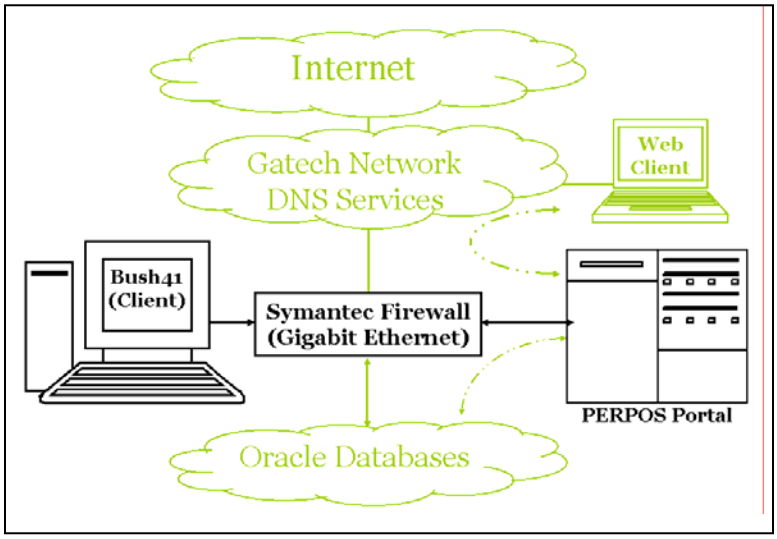

Figure 8. Test scenario 2 (local connection).

The scenario also consisted of 6 tests. Each was conducted under different network and security options to measure the performance of the portal in terms of its effective data transfer rates *(r)* as a function of (1) the operating mode of the portal *(m)*, (2) the cipher suite *(c)* used in a downloading session, and (3) the use of application inspection feature *(a)* of the Symantec firewall; *i.e., r=f (m, c, a).* The options used in the 6 tests are as follows:

Table 2. Local retrieval configurations.

Test #	m	c	a	Chart Label
1	unsecured	null	OFF	Bush41_unsecured_mode_Symantec - app_insp
2	unsecured	null	ON	Bush41_unsecured_mode_Symantec + app_insp
3	secured	DES-CBC-SHA	OFF	Bush41_DES-CBC-SHA_Symantec - app_insp
4	secured	DES-CBC-SHA	ON	Bush41_DES-CBC-SHA_Symantec + app_insp
5	secured	DES-CBC3-SHA	OFF	Bush41_DES-CBC3-SHA_Symantec - app_insp
6	secured	DES-CBC3-SHA	ON	Bush41_DES-CBC3-SHA_Symantec + app_insp

The results of all the six tests are displayed in figure 9 below. The chart indicates that the throughputs were the function of two main variables: the operating mode of the portal and the chosen cipher suite. The throughputs appeared to be unaffected by the use of the application inspection of the Symantec firewall. The results of tests 3 and 4 were indistinguishable on the chart, and so were those of tests 5 and 6 because the Symantec firewall was able to process incoming packets at a very high rate. The firewall software was running in a high-performance computer equipped with Gigabit Ethernet network interface cards.

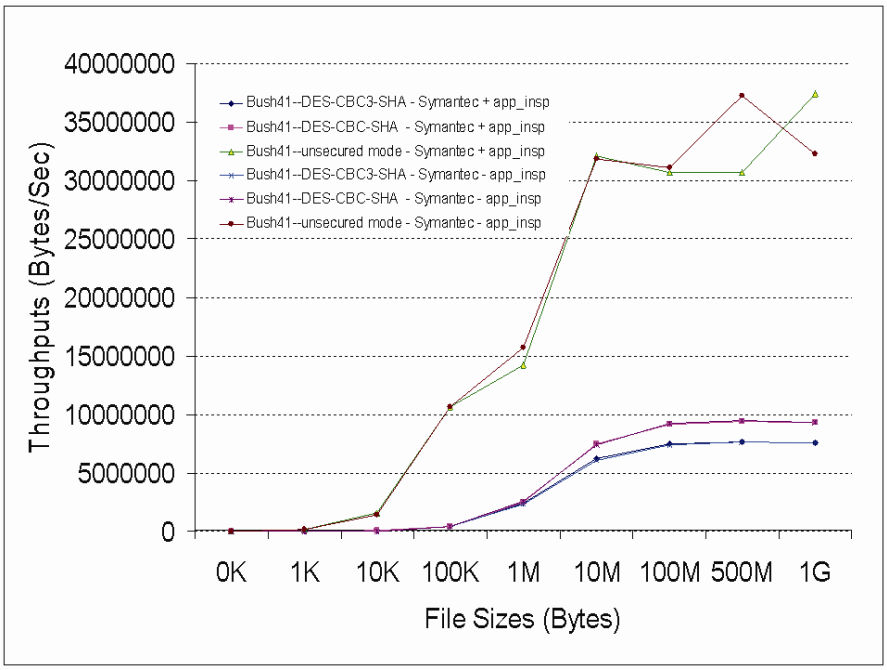

Figure 9. Comparative results of test scenario 2.

The results of the six tests in this scenario were categorized into two groups, corresponding to the dual operating modes of the portal: unsecured mode and secured mode. The average of each group was then computed and plotted in figure 10. The performance costs were then computed and plotted in figure 11. The costs were calculated based on the following formula:

$$\text{Performance Cost (\%)} = 100\ (T_i - T_0) / T_0$$

Where

T_0 is the average measured throughput in unsecured environment

T_i is the average measured throughput in secured environment

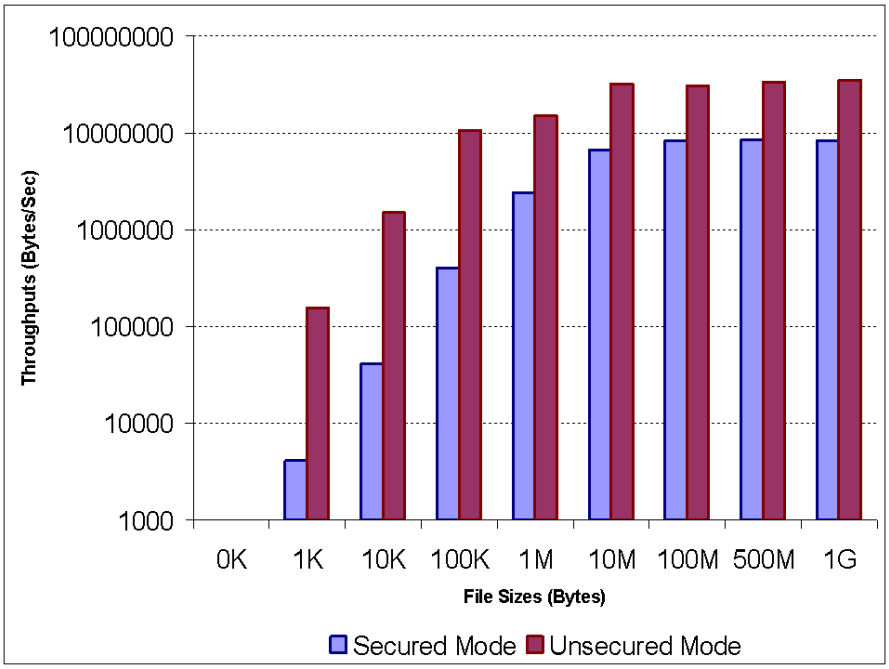

Figure 10. Average throughputs of test scenario 2.

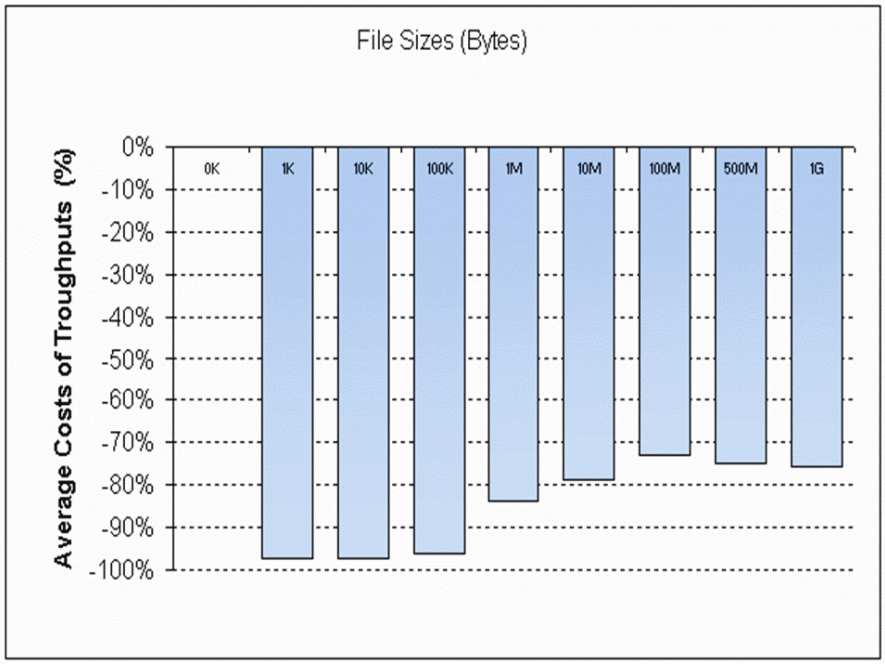

Figure 11. Comparative performance costs test scenario 2.

3. Conclusions and Recommendations

For the first time, it was possible to measure the performance of the portal operating under various configuration settings of the employed security products while conducting live experiments in the PERPOS test bed. The security products included the two recently acquired firewalls and the security services available at the portal. Initial experimentation with the portal uncovered the following facts about the test bed:

- The secured PERPOS portal provided the highest level of security services using the DES-CBC3-SHA cipher suite together with the version 3 of the secured socket layer (SSLv3) protocols. This cipher suite uses the RSA methods for key exchange and authentication of the portal, triple DES symmetric data encryption algorithm with 168-bit cryptographic key for data confidentiality, and the secure hash standard version 1 (SHA1) for data integrity.
- The SSLv3 protocols used in the portal are not FIPS-approved protocols, although they are widely used in the Web to secure the transfer of sensitive data over the Internet.
- The RSA public-key methods, the SHA1 secure hash algorithm, and the Triple DES encryption algorithm were government-approved standards for providing authentication, integrity, and confidentiality services, respectively.
- Although the DES encryption algorithms are government-approved symmetric data encryption algorithms, they are now considered obsolete and being replaced by the AES algorithms.
- A mix of Gigabit Ethernet (1000 Base-T) and Fast Ethernet (100 Base-T) network devices was used in the PERPOS laboratory. Fast Ethernet networks can support data-transfer rates up to 100 megabits per second (Mbps), whereas Gigabit Ethernet networks can support data-transfer rates up to 1000 megabits (gigabits) per second (Gbps). The portal systems had Gigabit Ethernet network interface cards (NICs), whereas the hardware platform on which the Nokia CheckPoint firewall ran was equipped with Fast Ethernet NICs. Moreover, the network switches and cables used in the test bed were designed mainly for the Fast Ethernet network. Whenever a slower network device was used in the path between the PERPOS portal and its client, the slower device determined the overall throughputs of the connection.
- The hardware platform on which the Symantec firewall ran was highly suitable for the high-performance, Gigabit network of the PERPOS portal.
- The hardware platform on which the Nokia CheckPoint firewall ran was designed for a small Fast Ethernet network—not for the high-performance, Gigabit PERPOS network.

Based on these findings, the following statements can be made about the experiments:

- The measured throughputs were considered preliminary and for internal uses only. They did not indicate or predict the actual capability and the capacity of the PERPOS portal environment.
- The quantitative results obtained from these experiments were used to assist the project team members in gaining further understanding of the functionality of the employed firewalls, their capabilities, and their impact on the overall system performance, especially on the data-transfer rates (effective throughputs).
- Experiences gained from performing the initial tests benefited (1) the integration and configuration of the portal with advanced networking and security technologies, (2) the identification of bottlenecks and deficiencies in the test bed, and (3) the development of performance measurement criteria and requirements.
- Initial results provided encouraging evidence that (1) the concept of defense in depth can also be implemented successfully at GTRI to secure and protect the portal of sensitive archives, (2) the evaluation method and its associated software tools could be re-used and enhanced to measure the performance overhead, and (3) the use of government-validated defensive security products at the portal would provide NARA with reasonable assurance that sensitive ERA can be secured and protected.

Below are recommendations for improving the PERPOS test bed environment and for enabling the PERPOS portal to rapidly respond to user requests for secured transfer of sensitive ERA over the Internet:

- Acquire, install, and configure government-approved cryptographic modules implementing the TLS protocols and the AES symmetric encryption algorithms. Successful integration of these modules will enable the PERPOS portal to offer the highest protection of sensitive electronic presidential records in public network using the cipher suite TLS_RSA_WITH_AES_256_CBC_SHA specification.
- Modernize relatively old and slow computing platforms, including notebooks, desktops, and hardware platforms on which security products run, e.g., the hardware platform on which the Nokia firewall operated.
- Upgrade network equipment, including switches, routers, cables, and network interface cards to support Gigabit Ethernet network technologies.

For the coming efforts, ARL recommends that the following tasks be conducted in FY06:

- Continue conducting empirical experiments evaluating the performance overhead induced by the deployment of defensive IA products at the actual PERPOS portal operating in unsecured and secured mode.

- Evaluate open-source, secured, virtual private network (VPN) products that use cryptographic tunneling protocols to provide authentication, confidentiality, and integrity services for ERA in public networks. The evaluation includes literature review and empirical experimentation.
- Continue assessing other types of government-validated IA products suitable for the protection and the empirical experimentation of the actual PERPOS web server. Products that need to be evaluated will include intrusion detection technologies and products, antiviral software, and security management tools.

4. Technical Contributors

Matthew Underwood and Jason Kau of GTRI and Son Nguyen of ARL were also technical contributors of this project. Mr. Matthew Underwood was the principal architect and administrator of the PERPOS computing test bed. Mr. Jason Kau was the network and firewall specialist. Mr. Son Nguyen selected appropriate government-validated firewalls and collaborated with Mr. Jason Kau to install and operate the firewalls. Mr. Son Nguyen also coordinated with Mr. Matthew Underwood to run the downloading scripts and to gather the results, which were analyzed and presented in this report.

5. Commonly Used Acronyms

DDR SDRAM	double-data-rate synchronous dynamic random access memory
DNS	Domain Name System
IP	Internet Protocol
HTTP	Hypertext Transfer Protocol
HTTPS	The secure version of HTTP
NIAP	National Information Assurance Partnership
RAID	Redundant array of independent disks
SSL	Secure Socket Layer
TLS	Transport Layer Security

6. Product and Vendor Information

- Operating Systems, Tools, and Utilities

Cygwin	http://www.Cygwin .com
Gentoo	http://www.gentoo.org
Microsoft	http://www.microsoft.com
Linux	http://www.linux.org
Unix	http://www.unix.org
Redhat	http://www.redhat.com

- Web Portal and Web Server Products

Oracle Application Server 10g	http://www.oracle.com
The Apache Software Foundation	http://www.apache.org
Transport Layer Security Protocols and Cipher Suites (OpenSSL)	http://www.openssl.org/docs/apps/ciphers.html

- Web Client Tools

wget	http://www.gnu.org/software/wget/wget.html
curl	http://curl.haxx.se

- Network Protocol Analyzers

ethereal	http://www.ethereal.com
ssldump	http://www.rtfm.com/ssldump

INTENTIONALLY LEFT BLANK.

Distribution List

ADMNSTR
DEFNS TECHL INFO CTR
ATTN DTIC-OCP (ELECTRONIC COPY)
8725 JOHN J KINGMAN RD STE 0944
FT BELVOIR VA 22060-6218

NATL ARCHIVES & RECORDS ADMIN
ELECT RECORDS ARCHIEVS PROG MGMT OFC
ATTN R CHADDUCK (2 COPIES)
8601 ADELPHI RD
COLLEGE PARK MD 20740-6001

GEORGIA TECH RESEARCH INSTITUTE
ATTN W UNDERWOOD (2 COPIES)
ATLANTA GA 30332

US ARMY RSRCH LAB
ATTN AMSRD-ARL-CI-OK-TP TECHL LIB
T LANDFRIED (2 COPIES)
ABERDEEN PROVING GROUND MD 21005-5066

US ARMY RSRCH LAB
ATTN AMSRD-ARL-CI-CN B NGUYEN (2 COPIES)
ATTN AMSRD-ARL-CI-CN G RACINE
ATTN AMSRD-ARL-CI-OK-T TECHL PUB (2 COPIES)
ATTN AMSRD-ARL-CI-OK-TL TECHL LIB (2 COPIES)
ATTN AMSRD-ARL-D J M MILLER
ATTN IMNE-ALC-IMS MAIL & RECORDS MGMT
ADELPHI MD 20783-1197

INTENTIONALLY LEFT BLANK